NERSC
A Computational Facility for the DOE Office of Science

Horst D. Simon

Director, NERSC Divsion

Lawrence Berkeley National Laboratory

Outline

- **NERSC Overview**

- **Assuring scientific quality**
  - **Advisory structure**
  - **Review process**

- **Maintaining cutting edge technology**

NERSC Center Overview

- Funded by DOE, annual budget $28M, about 65 staff
- Supports open, unclassified, basic research
- Located in the hills next to University of California, Berkeley campus
- close collaborations between university and NERSC in computer science and computational science
- close collaboration with about 125 scientists in the Computational Research Division at LBNL

National Energy Research Scientific Computing Center

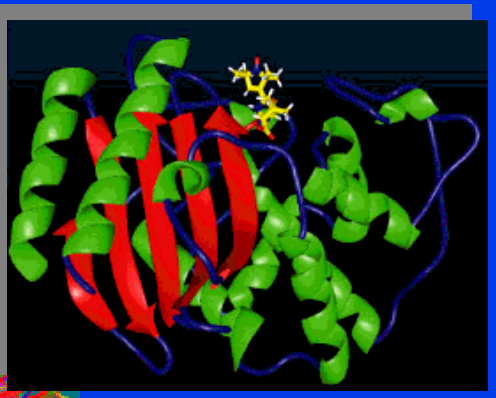

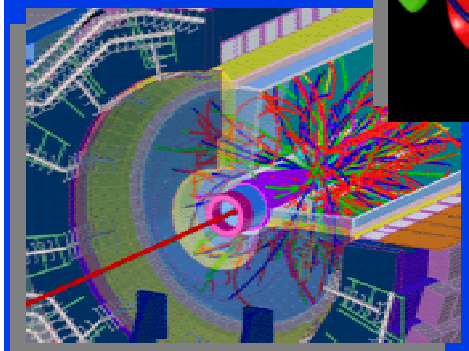

- Serves all disciplines of the DOE Office of Science
- ~2000 Users in ~400 projects
- Focus on large-scale computing

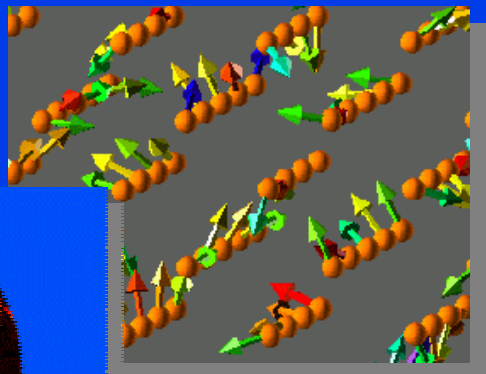

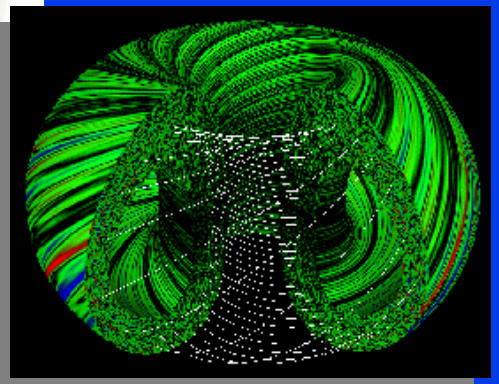

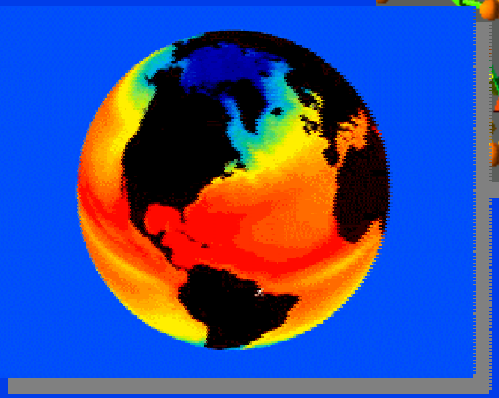

Upgraded NERSC 3E Characteristics

- The upgraded NERSC 3E system has
  - 416 16-way Power 3+ nodes with each CPU at 1.5 Gflop/s
    - 380 for computation
  - 6,656 CPUs – 6,080 for computation
  - Total Peak Performance of 10 Teraflop/s
  - Total Aggregate Memory is 7.8 TB
  - Total GPFS disk will be 44 TB
    - Local system disk is an additional 15 TB
  - Combined SSP-2 is greater than 1.238 Tflop/s
  - NERSC 3E is in full production as of March 1,2003
    - nodes arrived in the first two weeks of November
    - Acceptance end of December 2002
    - 30-day availability test near completed Feb. 2003
    - In full production March 1, 2003

TOP500 List of Most Powerful Computers

Rank	Manufacturer	Computer	R_{max} [TF/s]	Installation Site	Country	Year	Area of Installation	# Proc
1	NEC	Earth-Simulator	35.86	Earth Simulator Center	Japan	2002	Research	5120
2	HP	ASCI Q, AlphaServer SC	13.88	Los Alamos National Laboratory	USA	2002	Research	8192
3	Linux Networx/ Quadrics	MCR Cluster	7.63	Lawrence Livermore National Laboratory	USA	2002	Research	2304
4	IBM	ASCI White SP Power3	7.3	Lawrence Livermore National Laboratory	USA	2000	Research	8192
5	IBM	Seaborg SP Power 3	7.3	NERSC Lawrence Berkeley Nat. Lab.	USA	2002	Research	6656
6	IBM/Quadrics	xSeries Cluster Xeon 2.4 GHz	6.59	Lawrence Livermore National Laboratory	USA	2003	Research	1920
7	Fujitsu	PRIMEPOWER HPC2500	5.41	National Aerospace Laboratory of Japan	Japan	2002	Research	2304
8	HP	rx2600 Itanium2 Cluster Qadrics	4.88	Pacific Northwest National Laboratory	USA	2003	Research	1536
9	HP	AlphaServer SC ES45 1 GHz	4.46	Pittsburgh Supercomputing Center	USA	2001	Academic	3016
10	HP	AlphaServer SC ES45 1 GHz	3.98	Commissariat a l'Energie Atomique (CEA)	France	2001	Research	2560

FY02 Usage by Scientific Discipline

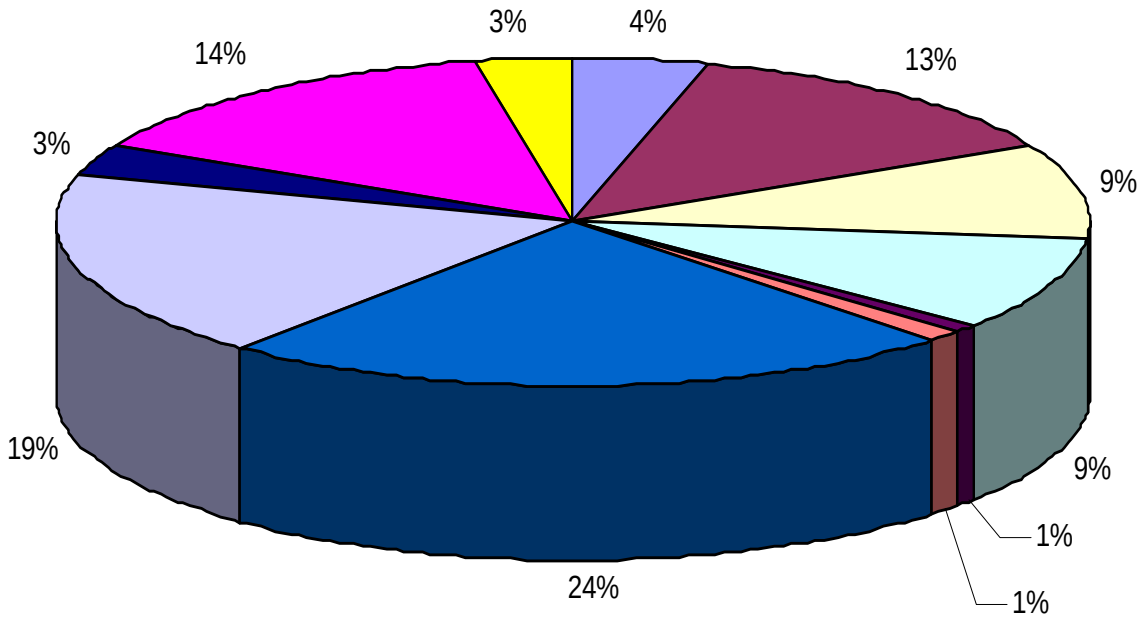

FY02 Usage by Institution Type

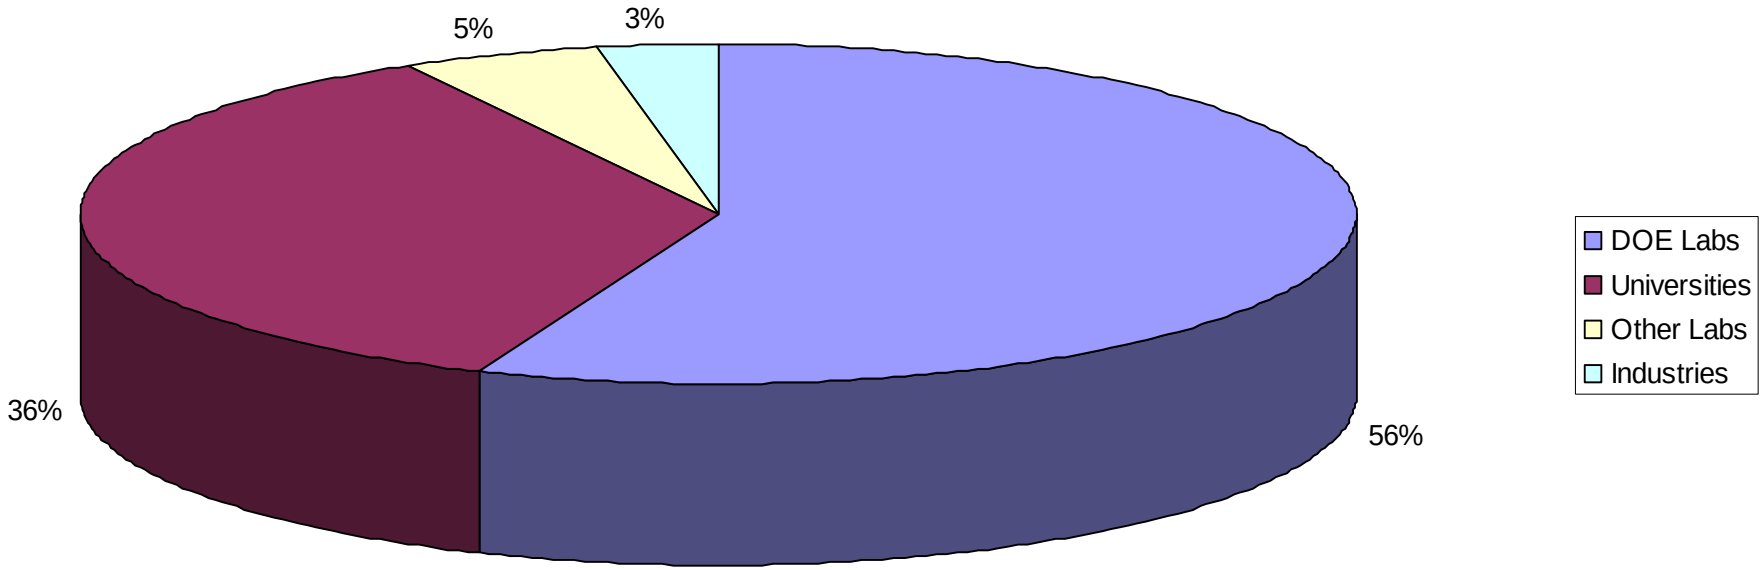

Outline

- NERSC Overview
- Assuring scientific quality
  - Advisory structure
  - Review process
- Maintaining cutting edge technology

Advisory Structure – NERSC Policy Board

- NERSC Policy Board (NPB)
  - Widely respected leaders in science, computing and technology (8 members)
  - Reports to and appointed by LBNL Director
  - Executive level advice on quality of science, long range planning, strategic issues
  - Meets once a year at LBNL

Advisory Structure
NERSC User Group

- NERSC User Group
  - All PIs with allocations are member
  - Members elect chair and executive board (21 members)
  - Reports to NERSC Director
  - Provides advice and feedback on current state of NERSC resources and their delivery to users
  - Develops requirements document (NERSC Greenbook)
  - Advocates NERSC within the DOE-SC community
  - Meets 1-2 times per year at user sites or national meetings
  - monthly teleconferences of NERSC management with executive board

NERSC Allocations Review Structure

- Request for allocations at NERSC are submitted annually through the ERCAP process, a web based proposal interface
- There is no science review, since the projects are peer reviewed in their programs at DOE

Review Panels for Allocations:

- CORP - Computational Review Panel
  - Scientific peer review of computational readiness
  - About 50 – 60 reviewers evaluating computational science, math., computer science, feasibility
- SAC – Scientific Allocations Committee
  - Program managers at DOE-SC take the CORP review as input and make allocation decision
  - Decide about DOE mission relevance for non DOE funded proposals

NERSC Allocations Review Structure (cont.)

- INCITE program
  - 10% of NERSC resources will go to a very small number of projects (2 to 3), which promise big impact in science
  - Open to all
  - Peer review based on quality of science AND computational feasibility
- Start-up allocations
  - Discretionary allocations made by NERSC Director (up to 5% of resources)
  - Help new users and projects to get started
  - Expectation is that after maximally 18 months the PI will apply for regular funding

Outline

- NERSC Overview
- Assuring scientific quality
  - Advisory structure
  - Review process
- Maintaining cutting edge science and technology

DOE Review: Strategic Plan

- Comprehensive peer review of the facility every 5 years
- Peer reviewed by 14 outside reviewers in late 2001
- Defines NERSC as general purpose, full service, capability center for FY2002 – FY2006

NERSC's Service Architecture

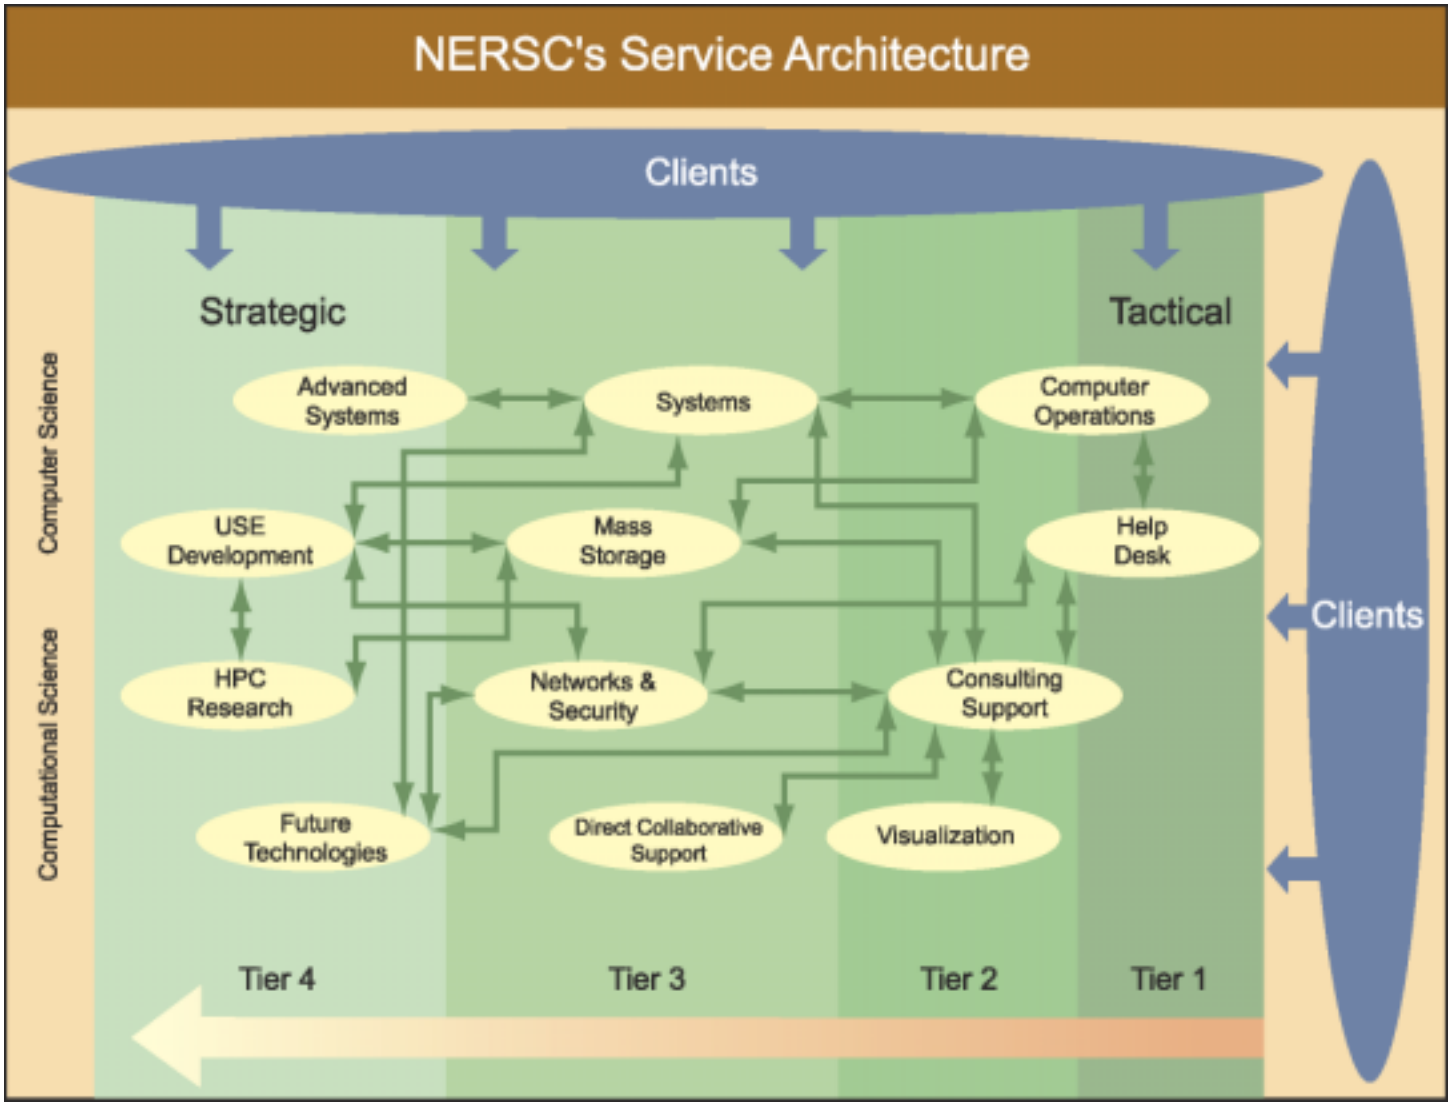

Advanced Development and Research Collaborations at NERSC

As a state-of-the-art facility NERSC must maintain a high level of continued expertise in computational technology.

- NERSC dedicates about 15 FTEs (out of 65 FTEs) to advanced development and research collaborations
  - For computer technology assessment and integration of new technology (7 FTEs)
  - For collaborations as computational/visualization experts as peers to the user community (8 FTEs)
- In addition 9 FTEs are user consultants, often working directly with users on research level issues

Computational Research Division (CRD) at LBNL

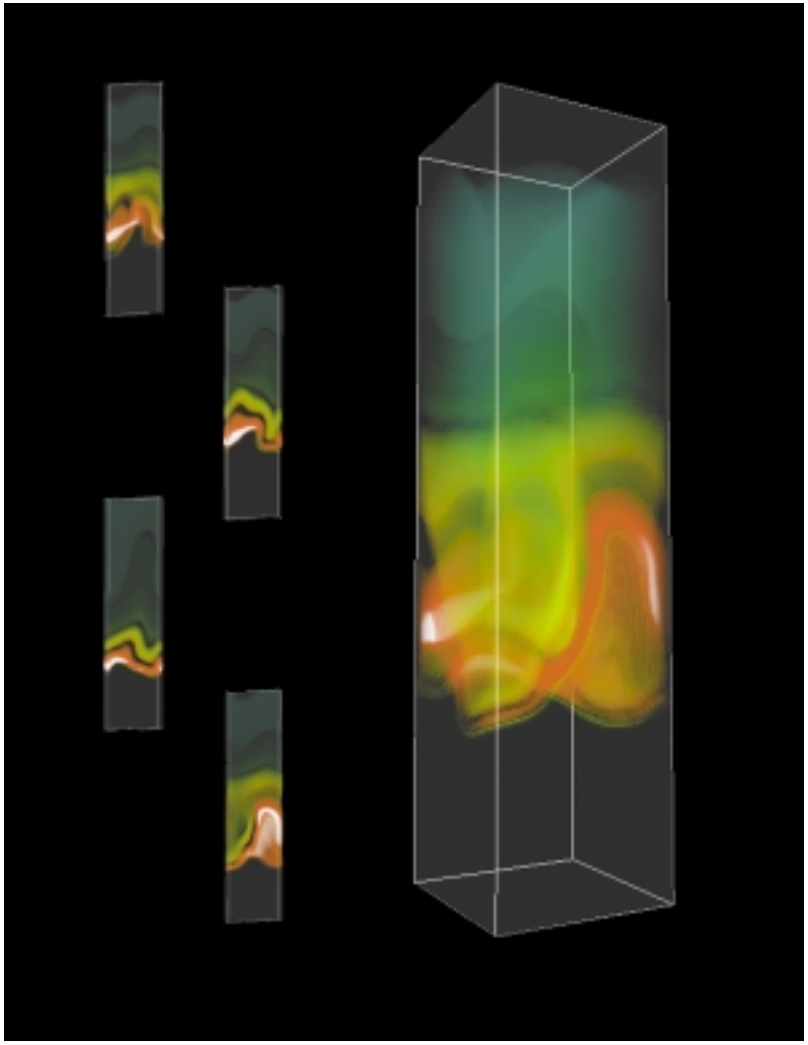

Horst Simon, Division Director

The Computational Research Division (CRD) creates computational tools and techniques that enable scientific breakthroughs, by conducting applied research and development in computer science, computational science, and applied mathematics. CRD consists of two departments:

1. High Performance Computing Research (108 FTE)
2. Distributed Systems (25 FTE)

- About $5M of SciDAC funded projects

Computational Research Division at LBNL and the NERSC Division

- A organizational balance between production facility and research activities is essential for the success of both:
  - NERSC benefits directly from the DOE funded math and computer science research projects
  - the requirements of NERSC stimulate computer science research and keep it closely relevant to the DOE mission
- NERSC has been a successful catalyst for technology transfer from DOE funded math and CS research in CRD into computational applications, e.g.
  - data management
  - distributed computing
  - CFD software

NERSC as a Division at LBNL

- As an integral part of an Office of Science multipurpose lab NERSC is well connected to the mission objectives of the Office of Science on a daily basis

- NERSC at LBNL has taken advantage of the rich intellectual environment at the lab and built new collaborations to the advantage of its national community

Computational LDRD Projects at LBNL

- In order to stimulate computational science in collaboration with, LBNL directed a significant portion of its LDRD budget into computational projects
- From 1996 – 2001 about 40 projects were started in collaboration with NERSC
- These projects had significant scientific results ***AND*** increased NERSC's national impact

LBNL Impact on NERSC

- Saul Perlmutter (Supernova Cosmology Project, LBNL) was co-recipient of Science magazine's "Breakthrough of the year award" (1998): universe expands at an accelerating rate.
- This project is a "graduate" of the computational science program at LBNL (since 1996); LDRD funded; started with the arrival of NERSC
- Simulations with radiation transport code and data analysis done at NERSC
- Continued collaboration with NERSC in SNAP project

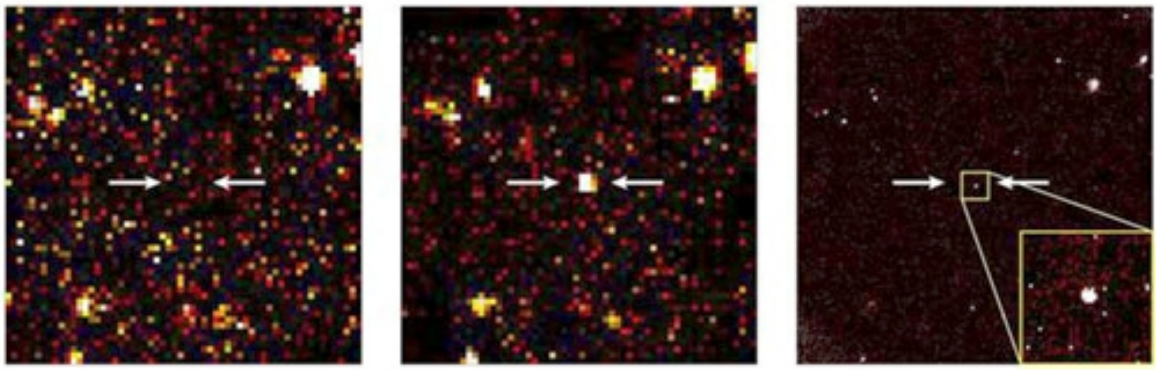

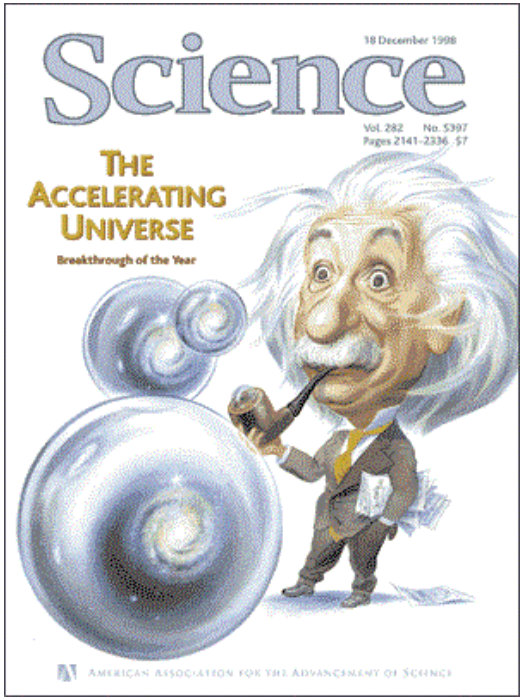

Impact of UC Berkeley and local community

- As an open facility in close proximity of a major research university NERSC is taking advantage of the local intellectual environment
- NERSC at LBNL has taken leveraged connections to UC Berkeley, EECS and other local academic institutions to the advantage of the DOE mission:
  - rapid introduction of new technology
  - leverage of computer science research
  - staff development and outreach
  - joint research projects

Summary

- NERSC has established a complex system of external advisory committees, assuring continued scientific quality and technical relevance of the facility

- NERSC has taken maximal advantage of local resources (LBNL, UCB, internal and CRD research) to advance the state of the facility and to increase its value to DOE